Frank L. Greitzer and Thomas A. Ferryman

Predicting Remaining Life of Mechanical Systems

ABSTRACT

A diverse set of challenges faces government and industry for life-cycle maintenance of aging equipment. Advancements in sensor and computer technologies are making it feasible to install sensors and small powerful computers on complex equipment to monitor the general condition or state (i.e., health) of that equipment. Methods for analyzing system status and health and for predicting system life expectancy need to be made more powerful, insightful, reliable, and robust for data collected onboard systems in real time. This paper reports on a project that is developing robust analytic methods for predicting remaining life of mechanical systems. The project is focused specifically on the investigation of a generalized statistical method for characterizing and predicting system degradation.

INTRODUCTION

A pervasive problem in both government and industry is the need to extend the useful life of systems. Economic pressures to maintain aging military and commercial equipment and vehicles (on the ground, on/under the sea and in the air) are very real. Even with relatively new equipment, there is a tremendous benefit of extending the time between overhauls and maintenance, reducing the probability of a failure in the field, and increasing appropriate preventive repairs. Major predictors of the need for maintenance are the type of use and operating conditions—such as environmental factors, duty factors, and service history—experienced by the product. Key to extending the useful life of each of these systems is the capability to extract information from the operational experience of this specific system to produce reliable diagnostics and prognostics about the state of the specific system and its remaining useful life. This capability may (and often does) require recording system data and analyzing the data using mathematical or statistical tools.

Prognostics are the process of predicting the future state of a system. Prognostics systems comprise sensors, a data acquisition system, and microprocessor-based software to perform sensor fusion, analysis, and reporting/interpreting of results with little or no human intervention, in real time or near real time. Many experts expect effective prognostics will result in reduced numbers and severity of failures (especially failures in the field), optimizing operational performance, extending the time between needed maintenance activities and reducing life-cycle costs. Implementing prognostics is a challenging task on several levels: 1) hardware and sensor technologies, 2) analytically effective predictive methods, and 3) organizational changes to capture the operational, maintenance and logistical benefits made possible by effective prognostic information. Still, the benefits to be gained by using effective prognostics dwarf the costs of the task.

Research and development under way at the Pacific Northwest National Laboratory (PNNL) is aimed at multilevel challenges of developing new sensor technologies for acquiring data, developing effective analytic methods for predicting system status using platform-based, real-time sensor data, and defining organizational concepts to help institutions make the most effective use of the prognostics outputs. One of the major drivers of this research program is the Life Extension Analysis and Prognostics (LEAP) project, an internally funded Laboratory Directed Research and Development activity (Greitzer, 1999). This paper describes the LEAP project and its focus on analytic methods to enhance the quality of predictions by prognostics systems.

BACKGROUND

Traditional maintenance practice either is a function of a somewhat arbitrary rule-of-thumb (e.g. perform maintenance every 90 days or 1,000 hours of operation) or it is reactive (i.e., performed when the equipment breaks). In the latter case, the necessary parts, supplies, personnel, and tools required for maintenance may not be adequately placed in the pipeline until maintenance is scheduled. As a result, when the actual demand is greater than expected and parts begin to fail early, the ramp-up time for maintenance may be steep and inefficient, particularly when personnel, parts, materials, or other resources are not readily available. The non-availability of resources can result in operational forces degradation, a large production deficit and substantially increased delays and costs. On the other hand, when the actual demand is less than expected, a traditional maintenance and acquisition system may suffer unnecessary costs associated with idle personnel and prematurely purchased parts.

Performing scheduled maintenance based on engineering judgment or mean-time-between-failure (MTBF) statistics is an attempt to reduce this problem. However, this practice typically results in equipment being serviced or replaced before doing so is necessary. Even worse, entire equipment replacements may be necessary when equipment fails before the expected time. When a schedule is based on an average (in this case, MTBF), one might expect it will specify too-frequent maintenance about half of the time and too–infrequent maintenance about half of the time. As a consequence, almost all systems will receive service or repair activity either too soon or too late (although it is possible that a few of the systems will be in need of service or equipment replacement at the very moment the schedule would indicate all systems need the service or equipment replacement). In addition, manufacturers generally follow the adage "an ounce of prevention is worth a pound of cure" and schedule maintenance extra early to minimize the probability of failures, which incurs unnecessary costs and further de-optimizes the process.

In contrast, prognostics can enhance the process of scheduling maintenance, ordering parts, and using resources. Figure 1 (Greitzer et al. 1999) shows the uses of onboard (platform) data in establishing wear databases and models and in life-cycle support. The left side of Figure 1 depicts the *non-real-time* activity of establishing wear models. Although many months of data from many pieces of equipment may be required to develop accurate, comprehensive, and stable wear models; some benefits begin to be realized immediately. Once these models are developed, they are downloaded to each piece of equipment for use in the real-time prognostics module. The right side of Figure 1 illustrates the use of *real-time* prognostics. Onboard the system, data from the platform's computer processor(s), sensors, and the environment are used in conjunction with the wear models to predict performance and wear in real time. Once a failure or excessive degradation is predicted, data about the impending event could be forwarded to a central logistics system. The data forwarding could be done either real-time, if the conditions warranted, or via periodic downloads. In addition, if conditions warranted, the operator could be informed real-time. This vision enables maintenance to be scheduled based on these data, and necessary equipment and parts are ordered to arrive just in time for the maintenance. The operator may choose to modify the maintenance routine based on predicted health.

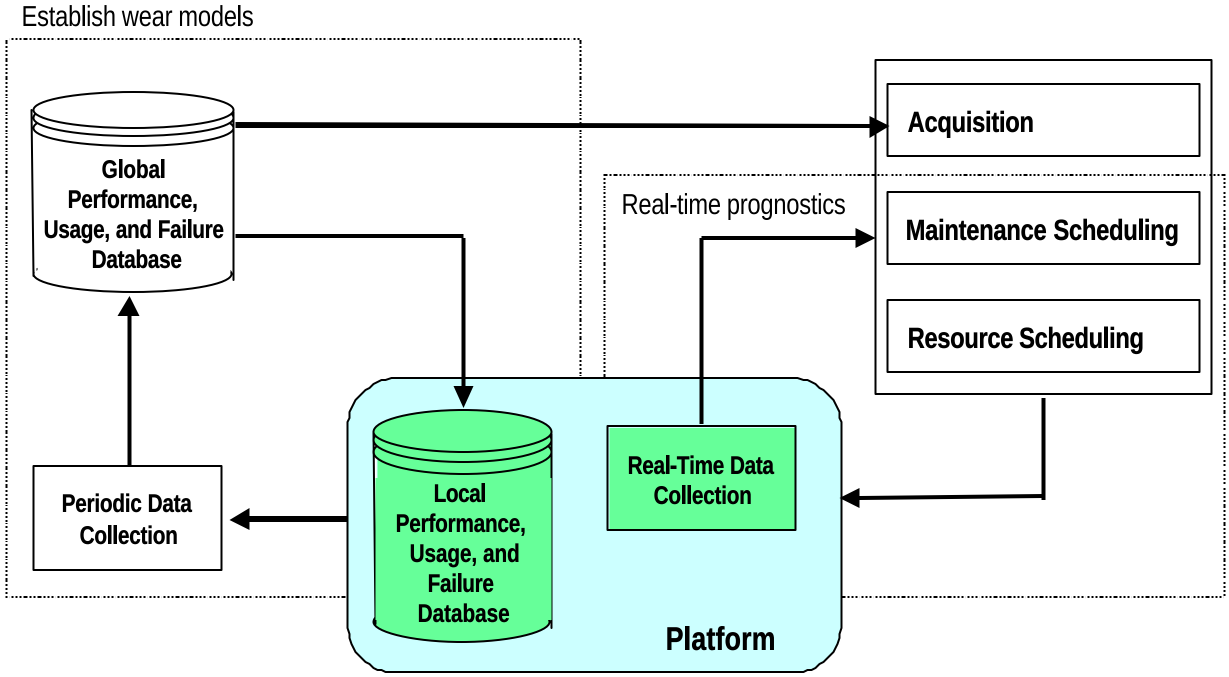

Figure 1. Prognostics-Enabled Scheduling

Two types of onboard health-monitoring data are distinguished: 1) information outputs comprising a relatively small amount of coded data indicating health status of specific components and systems; and 2) raw data and detailed processed data used by or developed from onboard analysis software. Health monitoring outputs are used throughout the prognostics health monitoring environment, from operators to maintenance personnel and on up the logistics/maintenance hierarchy. However, the higher up the hierarchy, the more aggregated the data become. At the highest levels, statistical summaries and trends are used to support decisions such as procurements and life-cycle analyses. At the lowest levels, platform-specific data are used to support maintenance operations. Raw/processed data that have been captured on board and stored in mass storage (e.g., hard disk drive or flash memory) may be collected periodically and used to improve capabilities and accuracy of the prognostics software. For this reason, raw/processed data are collected for development/research purposes. However, off board analysis software can be provided to the field mechanic; this software can "replay" the raw data stored on a platform to support troubleshooting.

Data collected during field operations may be used to develop and refine prognostic algorithms that facilitate proactive life cycle and maintenance activities. Failures must be predicted early enough to enable maintenance and acquisition systems to prepare for the workload and required replacement parts ordered, significantly reducing maintenance time. The ability to maximize the benefits from prognostics will require proactive business practices. Organizations will achieve these benefits by evolving beyond their reactive processes and adopting new proactive objectives. Information to support decision-making must be directed at different levels: for the operator (What needs to be done? And when?), the maintenance technician (What needs to be repaired? And when? What parts need to be ordered? And when?), the engineer (What system weaknesses are apparent from the field usage?), and the accountant and life-cycle manager (What condition-based maintenance practices can result in minimizing costs? What systems are anticipated to be needed over the next time period, and how can they be acquired at minimal cost?). Effective asset management results from integration of these perspectives.

Figure 2 illustrates a high-level concept for a prognostics-based logistics/maintenance infrastructure architecture, based on ideas described elsewhere by the authors (Greitzer et al. 1999). This type of logistics/maintenance architecture would be applicable to a company or organization that employs large fleets of vehicles that need to be kept in service. Not all aspects need to be implemented to begin to achieve benefits, but maximal benefits will come after each aspect is implemented, integrated, and refined.

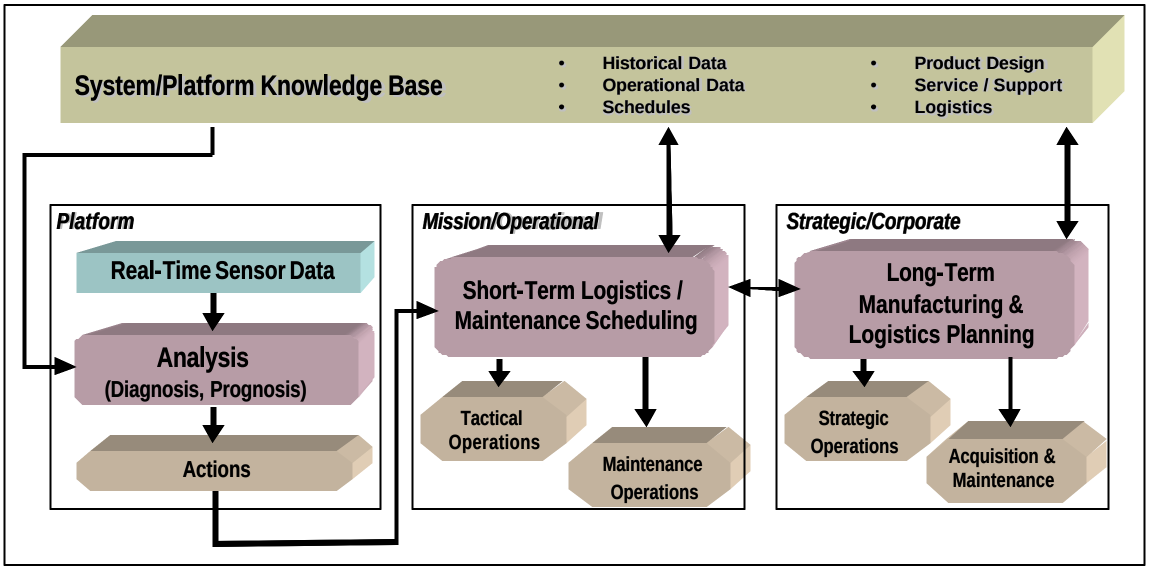

Figure 2. Logistics Structures Enabled by Prognostics-Based Health Monitoring.

APPROACH

To predict a *failure*—the inability or at least serious degradation of the platform to perform its intended function—three things typically must be known:

1. The system’s current degree of fault as quantified by a Figure of Merit (FOM)
2. A theory about the progression of the fault, so as to postulate the system’s degree of fault at a particular point in time in the future
3. The level of the fault, as quantified by the FOM, that will produce a *failure* of the platform.

The specification of these factors, which is necessary to perform diagnostics and prognostics, is typically done through engineering/analytical studies such as Failure Modes and Effects Analysis (FMEA). These analyses and expert judgments yield descriptions of how the system fails, what faults can be measured given available sensors, and the values expected for these sensors when these failures occur. These analyses and judgments are based on theoretical/physics-based determinations, archived manufacturer’s data, historical field data, and real-time data collected during operation of the system in the field Except in unusually simple cases, the FOM determined for a fault, failure, or system condition is a function of a combination of sensor values (i.e., through sensor fusion), rather than representing a single sensor. As a result, it is typically *not* sufficient to monitor and trend individual sensor values, independently, to perform diagnostics and prognostics. Thus, although the statistical methods discussed in this paper apply to single sensors, the following discussion uses the more general (and more often appropriate) FOM as the quantity to be monitored.

The ability to predict, in real time, the future state of a system based upon sensor data collected from an operating system requires analytic methods that must overcome inherent problems with dynamic data—namely, *dynamic variation* and *independent variable* selection.

Dynamic variation refers to the fact that the FOM will have some uncertainty in its estimated value under all conditions. FOMs are generally mathematical functions of one or more variables that have uncertainty in their measurements. Normal variation in the FOM and underlying sensor values must be distinguished from values that indicate degradation. An inevitable trade-off exists between using a large set of data to reduce the consequence of noisy sensor values and inherent system variability and using a smaller set of data to be responsive to changing system characteristics that may occur when a system health problem begins to manifest itself.

A second challenge relates to the selection of *independent variables* to be used for prediction. Prediction, by definition, implies the estimation of a parameter at some future point in time. Here, the use of the term "time" may be misleading, as it is clear that elapsed time or calendar time is a poor unit of measure for at least some of the failure mechanisms of a platform. Alternate and potentially better manifestations of the independent variable *time* might be *running time*, *cycles*, a measure of *work produced* (e.g., joules or torque-time), or any of dozens of others.

The approach employed in the LEAP research was to identify and investigate different statistical and analytic methods for improving the ability to predict future states of a system. Candidate statistical methods include multivariate regression, Bayesian regression, time-series analysis, and discrimination or clustering analysis. Analysis may focus on a single parameter or multiple parameters. For single-parameter prognostics, statistical analyses may be performed simultaneously on each real-time data source. As data are collected, regression models are applied to the data to determine trends in FOMs. These FOMs are compared, in real time, to metric failure limits that are established offline. The point of predicted failure is calculated as the intersection of these two lines (see Figure 3). Uncertainty intervals (dashed lines surrounding the trend lines) also may also be derived to incorporate uncertainty estimates into the prediction. In Figure 3, predicted time of failure is indicated by time t_2. The method estimates failure occurrence unlikely before time t_1 and unlikely to occur after time t_3.

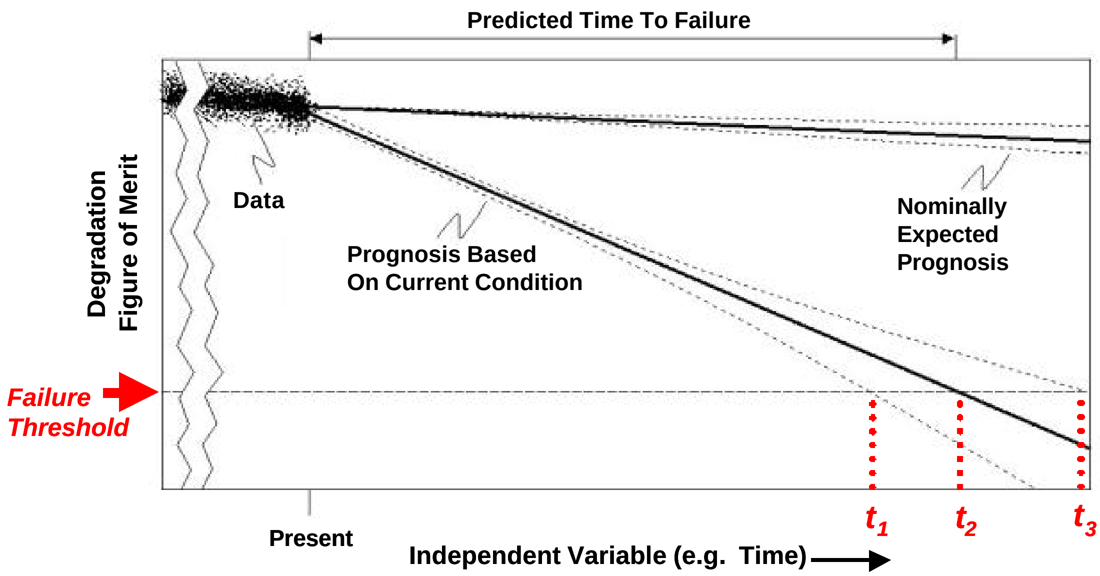

Figure 3. Regression Lines Intersecting Failure Threshold Indicate Predicted Time to Failure.

The amount of data included in the analysis affects the prediction. Use of large amounts of data spanning a long window of data acquisition tends to yield more stable, less variable predictions. However, it also may yield a prediction that is less sensitive to recent changes (see upper regression line labeled "Nominally Expected Prognosis" in Figure 3). Use of a smaller data set spanning the most recent operating history tends to produce predictions with more variability but more sensitivity to current operating conditions (see regression line labeled "Prognosis Based on Current Condition" in Figure 3). The goal of the LEAP project's prognostics analyses is to choose among varying window sizes to maximize the system's adaptability to change while maintaining an acceptable amount of predictive variability/uncertainty.

KEY TECHNICAL CHALLENGES

Given the nature of real-time diagnostics and prognostics, a number of factors drive the consideration of the mathematical/statistical algorithms that can be used practically in an analysis method. The algorithms must be selected appropriately for the available data, processing and storage capability, and analysis needs. Power requirements, heat generation, ruggedness of the diagnostics/prognostics system (CPU, dedicated sensors, and other components) must be considered. Naturally, cost and availability are major factors. The following discussion sources on key challenges regarding data sources, platform characteristics, uses of the results, and predictive analysis methods.

Data Sources

In addition to using the current and past condition of the platform to diagnose and prognosticate the health of the system, one should incorporate insight gathered from a variety of other sources. Manufacturer's data, which sets nominal periodic maintenance periods, might be used as input to a previous distribution regarding the system health, but it would be better to use the manufacturer's engineering data, which is based on analysis, computer models, and some related field experience. Field experience from other similar or related systems could also be reflected in a previous distribution. These data should be collected throughout the lifetime of the set of similar platforms and used to update what might be called the *nominally expected performance baseline*. The data updates would be accomplished by downloading the latest prognostic routines from a central analysis site database (possibly via telemetry or other automated means). The central location processing will not need to be real-time and can involve input from multiple platforms. The onboard algorithms would need to be designed to expect this information so as to not require a change in the programs.

Key data variables can be accessed from the platform's CPU by connecting to the platform's data network. The platform may have an embedded CPU that monitors the condition of the platform continuously. These data likely would be very useful for diagnostics and prognostics. Additional sensor data, such as oil temperature, air pressures or differentials, or vibrations at select locations, may be useful. These data are likely to be noisy and may have time-dependencies, which might include anything from low-frequency sinusoidal periods to high-frequency correlated noise. The noise may depend on the mode of the platform. In addition, the noise itself may be indicative of the underlying problem, so smoothing out the noise may be counterproductive.

Some data may come from dedicated sensors or sensors that are independent of the CPU. Some of the sensor data (whether acquired via the host CPU or dedicated sensors) may reflect *bad* data due to broken sensors or wiring, or miscalibration, electromagnetic interference, and other logistical problems. Unless this problem is mission-critical to the platform, it will not be fixed immediately. Consequently, diagnostics/prognostics efforts must continue to perform effectively, albeit with possible graceful degradation. The transition of the data quality from *good* to *bad* may be indicative of the platform experiencing a problem and if so, should affect the diagnosis.

Platform Characteristics
The platform may be complex and its nature as reflected in the data may be varied. The platform may have designated modes of operation such as *startup, warmup, run,* and *shutdown.* The data may disclose different modes of operation, such as, *initiation, startup, cold idle, warm idle, fast idle, low load, heavy load, turn-off,* and *cycle-down.* Platform design modes such as these are readily related to data modes. However, this may not always be the case, although it is desirable. The data analysis that enables effective diagnosis/prognosis may need to be related to the data modes. In addition, some platform problems will not be evident in some modes of operation (design or data), but they might be evident in other modes.

Platform behavior over time would often be expected to be stable or experiencing a very slow degradation. Failures might manifest themselves with a sudden shift in FOM, a sudden change in the rate of change in FOM (a slow degradation may become a rapid degradation), an accelerating rate of change of FOM, or any of several other possible manners. Ideally, a prognostics system will not only detect, quantify, and predict levels of performance of the steadily degrading system, but also will detect conditions that indicate a change in the probability of a sudden change in behavior.

The nature of what makes the platform degrade may be complex as well. Degradation may be related to calendar time (e.g., radiator hoses age as time passes), operation time (operation can stress metal to the point of fatigue), distance driven (e.g., tires), cycles (e.g., starters), fuel consumed (carburetor or fuel injector), work produced (in a physics sense, e.g., elevators), or any of numerous other factors not yet envisioned. Recognition of specific degradation factors may facilitate the formulation of the models used in the prognostics efforts.

Uses of the Results
The onboard diagnostics/prognostics system performs key functions associated with data collection, validation, algorithm processing, diagnostics processing, prognostic processing, and reporting. Appropriate information would be reported to different users according to their needs:

- Onboard operator needs
  - key conditions or constraints
  - imminent system problems
  - short-term prognosis
- Onboard maintenance leader needs
  - current operation characteristics, diagnostics with enhanced details
  - near-term prognosis with enhanced details
- operations scheduler needs
  - current diagnosis, near- and mid-term prognosis
- logistician needs
  - summary of diagnosis
  - near-, mid-, and far-term prognosis
  - historical operating characteristics, statistics, costs.

These four roles might be performed by one individual or four different organizations, depending on the platform and the organizational structure within which it operates. The level of detail of the reported information in each case will vary, but all will be based on the same core analysis, which obviously must support the highest level of detail.

Predictive Analysis Methods

Dozens of algorithms can be considered for use in a diagnostic/prognostic system, including but not limited to linear regression, linear multiple regression, time series analysis, Bayesian dynamic linear models (including Kalman filters), and non-linear regression and multiple regression. Also, multivariate versions of the aforementioned algorithms may be considered. The following should be considered in the evaluation of the appropriate algorithm to use:

- There always is a trade-off between false alarm rate and responsiveness to detect and report a change. Also, as pointed out in [ref], the system ideally will adjust the trade-off to reflect the costs and benefits of the system and the impact on the level of credibility that the user associates with the diagnostics/prognostics system.

- The data may be exceedingly noisy, especially when compared to 1) the potentially subtle trends that must be detected and 2) considerations when the projection is far into the future.

- The amount of data may never be sufficient to estimate a large number of model parameters. The platform may be relatively new when it is put to mission-critical use, and the prognostics need to be dependable. Some systems may have a relatively short life between major overhauls and never will acquire "large" amounts of data. The system may have minor maintenance/repairs that will affect the performance levels of some subsystems.

More specific concerns that relate to diagnostics/prognostics systems include:

- The system may change modes and conditions frequently, so relatively rapid re-calibration should be targeted.

- The system may have considerable related information, such as performance of similar systems or manufacturer's design specifications. To the extent practical, this should be used, but not to overwhelm the data from the individual system being monitored.

- The trends are likely to be very "gradual" (i.e., a small [or no] degrading slope).

- When a problem occurs, it is desirable to detect it and update the prognosis quickly.

- Too much flexibility in the model is likely to yield "good" fits that are an artifact of the noise rather than the true underlying signal. This can be aggravated further by the prediction into the relatively distant future. Allowing higher-order models provides greater fitting potential, but with extrapolation (as prognostics does, of course), the magnitude of a slight error in the coefficients gets magnified a great deal. Limiting the model to a first-order fit helps prevent the exaggeration of the extrapolated values associated with minor coefficient estimation errors.

- Use of all available data is desirable, as it will generally refine the estimate. However, if a platform undergoes a transition from one state (healthy) to another (broken or degrading), there is a need to detect the change rapidly. Use of a linear fit with a short window of data can be overly responsive to noise. Using a long-duration data window can be lethargic to change.

- Even the best predictions are invariably in error to some degree. Uncertainty intervals relate something about the accuracy of the predictions. They are critical in this application. Predictions that do not allow for the uncertainty of the interval will mislead the users and undermine the credibility of the system.

DEVELOPMENT AND ASSESSMENT OF PREDICTION METHOD

Initially, it was considered important to keep the model very simple because the data were limited and some extrapolation was necessary. Eventually, more sophisticated models could be developed to take more appropriate advantage of greater amounts of data representing prior information. Thus, this exploration was limited to a simple linear model, Y regressed on X, with a variety of independent variables being used as the X variable and a number of FOMs being used as the Y variable. This analysis was modified from a simple regression by incorporating a method to select among a small set of possible window lengths. The method developed to accomplish this selection can be compared to a variation of Bayesian Model Averaging.

The selected technique, named LEAP-Frog,[a] assumes the platform to be in a steady state of health for the most part. The goal is to predict future health without being overly sensitive to the noisy (and possibly correlated) data. At the same time, the technique must be responsive to changes in platform performance so a change in the platform's health can be detected and the predictions adjusted accordingly.

The prediction goal is to make a prediction at the *current time* for the value (and uncertainty intervals) of the FOM at a *future time* given all past data, and a relatively small set of models/time windows. The method begins with a regression analysis using large time window for data acquisition, which will likely give the best estimate of the regression fit, if the platform is at a constant rate of change of health (maybe steady with a slow rate of degradation). This prediction and an uncertainty distribution about the estimates are tested to see if it the prediction is reasonably compatible with the most recent data points. If so, then the regression is used. If not, then the size of the regression window is reduced and the analysis is repeated. This method continues until it yields a small enough window that is compatible with the most recent data points. As a result, this method can detect if the most recent data points indicate a change from the long-term regression (as would be the case if the platform had a change in FOM). In the end, the method uses the longest regression window that does not result in evidence (based on the most recent records) that refutes the assumption of a good linear fit; and then uses this window to predict the future state of the system or it's remaining useful life.

Another aspect of the method is to use different independent variables. Certainly, numerous regressors are possible, and these can be combined to construct a derived independent variable made up of a combination of multiple regressors that may result in better prognostics. At this stage of the development, however, the analysis is limited to regression on only one variable at a time, assessed within the context of the LEAP-Frog regression method. In future work, considerable improvement over these results is possible if more sophisticated combinations of independent variables are used.

To explore the method's effectiveness, it was used to analyze specially designed simulated data and actual field data from a U.S. Army Abrams tank. These tests are discussed below.

Tests With Simulated Data

To construct the simulated data analysis, a system was hypothesized with a FOM that started at 100 and degraded slowly downward. Two situations are considered: (A) slow steady degradation for 150,000 seconds (shown in Figure 4a), and (B) slow steady degradation for 80,000 seconds followed by a sudden change to a fast degradation for the remaining 70,000 seconds (shown in Figure 4b).

---

[a] Let us explain the derivation of the name Leap Frog. Pacific Northwest National Laboratory funded this work under the Lab-directed R&D program Life Extension Analysis and Prognostics (LEAP). The analysis technique can be described as a series of regressions which have the more effective regressions leap over the less effective regressions. Hence, LEAP-Frog.

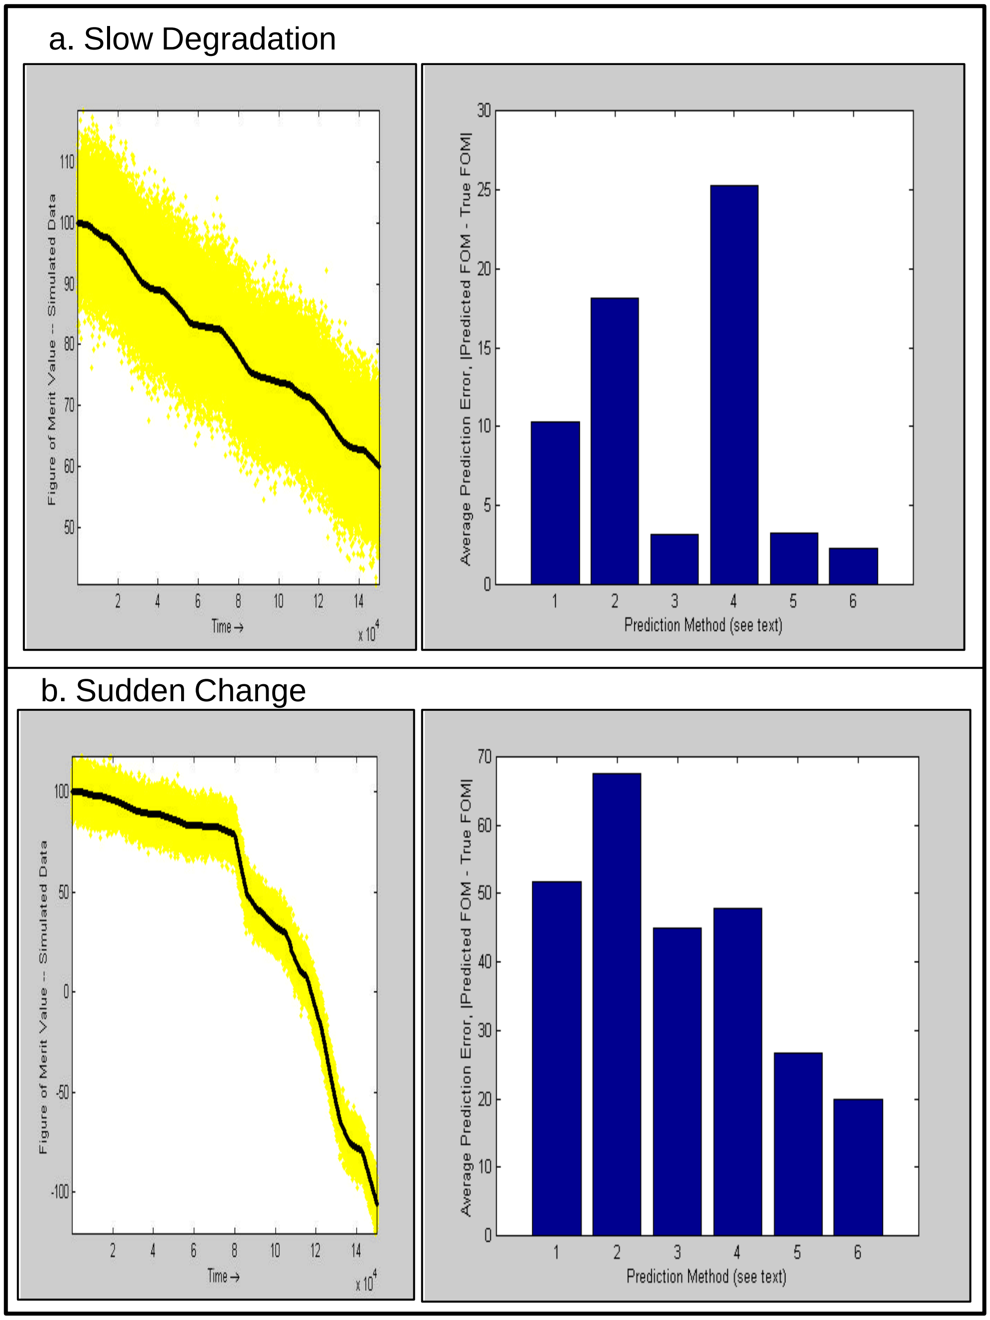

Figure 4. Comparison of Regression Methods for Predicting Performance of Simulated Data That Follow One of Two Degradation Patterns: (a) Slow Linear Degradation; (b) Slow Linear Degradation Followed by Sudden Change To Fast Linear Degradation.

The black-lined curves on the left sides of Figures 4a and 4b show the simulated true FOMs for Condition A and Condition B, respectively. The data points (plotted in dots) used for the simulation were additive N(0,5). Predictions were made every 500 seconds up to second 110,000. The predictions were made for 40,000 seconds in advance of the "current" time using six different prediction methods.

In the right sides of Figures 4a and 4b, four standard or typical prognostic methods are compared with two variations of the LEAP-Frog method. The method numbers indicated in the figure correspond to the six prognostic methods enumerated below, four of which are standard methods and two are based on the LEAP-Frog method. The four standard methods (respectively) predict future performance based upon:

(1) the last value of the FOM
(2) the average value of the FOM since the start of data collection
(3) the regression of the FOM on all the data since the start of data collection
(4) the regression of the FOM on the last 1,000 records.

The LEAP-Frog method provided two additional variations:
(5) a linear independent variable (such as time)
(6) a non-linear independent variable (such as distance driven).

Histograms on the right sides of Figures 4a and 4b show the average prediction errors for the six alternative regression methods under the two degradation patterns. For Condition A (slow linear degradation, shown in Figure 4a), it is not surprising that Method 3 (regression using all of the data) performs best, along with both LEAP-Frog methods (the latter methods would perform about the same as Method 3 because there would never be a reason to "jump" to a regression line with a smaller window). For Condition B (slow steady degradation followed by a sudden change to a fast degradation, shown in Figure 4b), the LEAP-Frog regression methods are best. This illustrates the ability of the LEAP-Frog method to adapt to a rapidly changing situation.

Gas Turbine Engine Field Data

In cooperation with the U.S. Army, several (to be consistent with the next sentence) M1 Abrams tanks were instrumented as part of a research and development program at PNNL funded by the US Army Logistics Integration Agency. In this project (called REDI-PRO, for Real time Engine Diagnostics-Prognostics), a prototype REDI-PRO system comprising sensors, data acquisition hardware, and analysis software was designed, developed, and installed on several Abrams tanks. Data were collected during routine operations of the tanks. One tank experienced a problem that was verifiable: an air filter clogged and caused the engine to shut down because a switch on the tank that otherwise would have prevented this event was inoperative. Approximately 7 hours of data were obtained from the instrumentation suite. To analyze the data, a FOM variable was formed by subtracting air pressure measured just *downstream* of the air filter from the air pressure measured *upstream* (just before that air filter). ("Before" and "after" are meant to relate the sense of normal airflow through this gas turbine engine.)

The available REDI-PRO data are shown in Figure 5. As shown in panels (a) and (b), in data collection up to twelve days before the air filter clog event, the REDI-PRO data acquisition system had compiled about 3.3 operating hours of data over about 3.5 months (the REDI-PRO data acquisition system was not acquiring data continuously). Based on the available data as of that time, the REDI-PRO/LEAP-Frog analysis shows no problem with the air filter (the trend line is fairly stable and flat, indicating no problem).

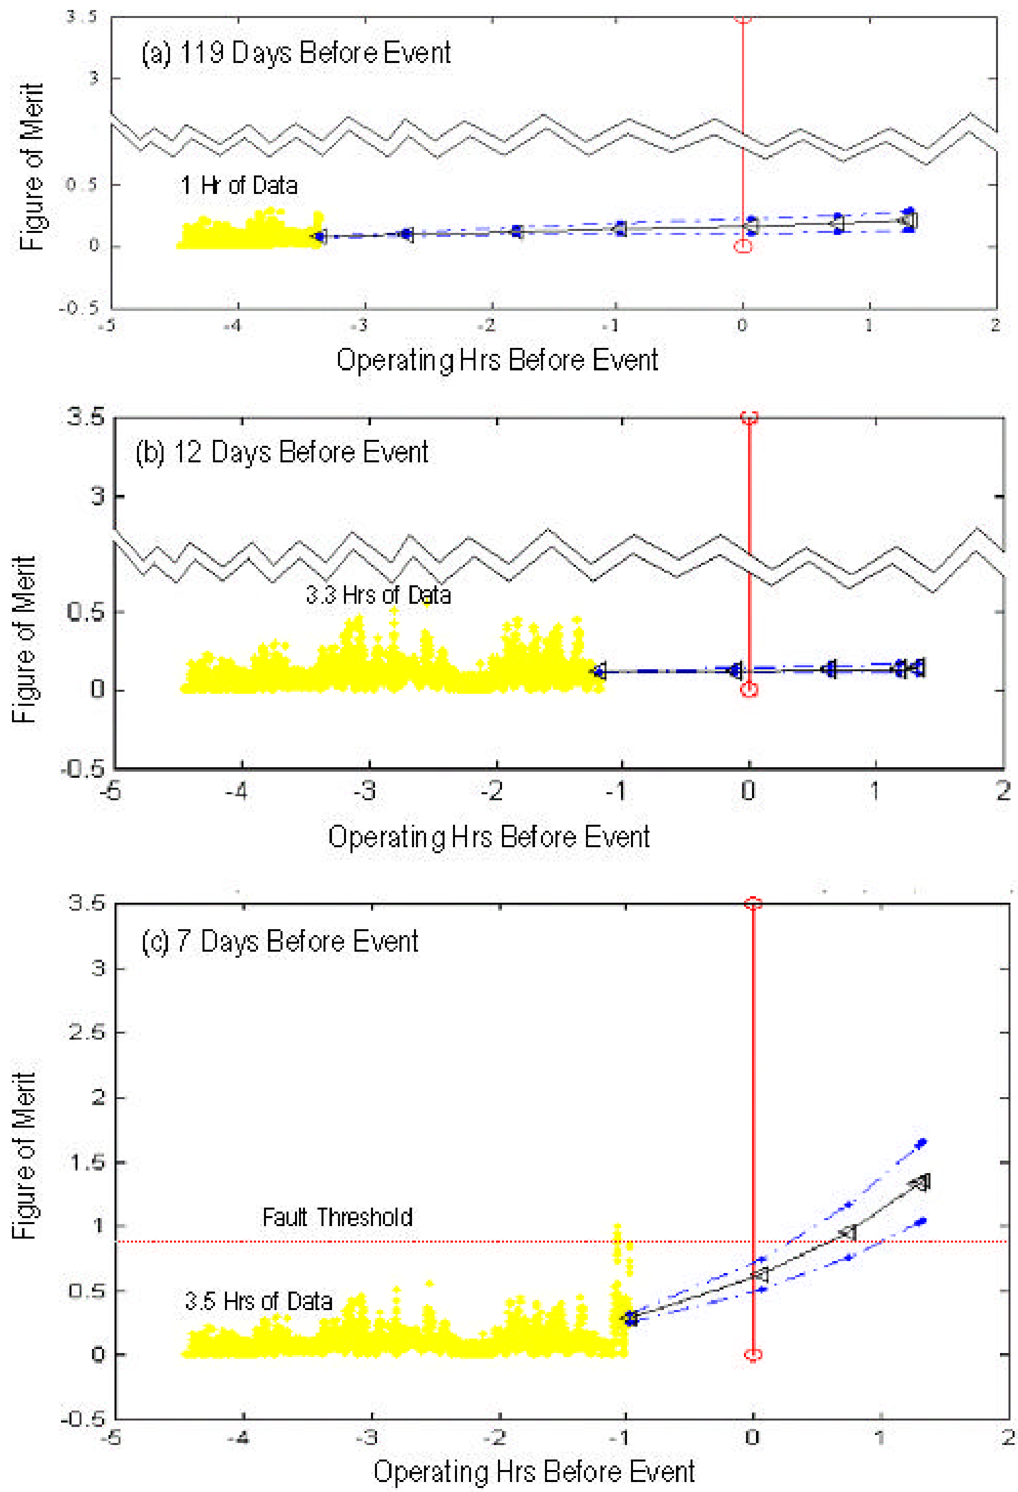

Figure 5. REDI-PRO Data Showing Prediction of Air Filter Clog Event

Figure 5c indicates that seven days before the event, the REDI-PRO system once again is operating. With 3.5 hours of data accumulated out of a total of about 43 operating hours, the LEAP-Frog regression analysis now indicates a change in engine life expectancy. The prediction is that the Air Filter Clogged light would illuminate (signaling a pending fault) after another 1.5-2 operating hours. The tank was not operated for the next five days, but it was run about 4 hours on the day before the event occurred (however, REDI-PRO captured a negligible amount of data because it was turned off most of the time). On the day of the event, an engine failure occurred while REDI-PRO was acquiring data. The Air Filter Clogged light had not turned on.

Subsequently, it was determined that the failure occurred because the Air Filter Clogged switch was inoperative (wired incorrectly), preventing the light from ever illuminating. Maintenance staff did not change the air filter because the light never came on. As a result, the air filter became so clogged that the engine failed to start. It is surmised that if the light had illuminated the day before the failure (as predicted by REDI-PRO), the filter would have been changed and the failure would not have occurred.

Because REDI-PRO was turned on and acquiring data less than 10% of the time the engine was operated, these predictions were made under less than ideal circumstances and should be taken as illustrative only. Continuing work on the REDI-PRO program with the US Army is aimed at obtaining additional field data, and data from test cells, that will help validate and refine the diagnostic algorithms used to generate FOMs for this engine. In addition, possible application of the REDI-PRO system and its LEAP-Frog predictive methods is being considered for upgrades or replacements to this engine.

NAVAL APPLICATIONS

Naval applications of the prognostic technology discussed in this paper are appropriate to wherever there are maintenance or logistical problems or potential problems. The REDI-PRO application demonstrated on one type of turbine engine is readily extrapolated to other types of turbine engines, to diesel engines, and other complex machinery. The Navy is already collecting operational data for some of its propulsion and power generator systems via the Integrated Condition Assessment System (ICAS) that is being administered by the Naval Surface Weapons Center, Carderock Division. Like the REDI-PRO system described in this paper, ICAS collects sensor data on various thermodynamic sensors at a rate of once per second. If existing, fielded sensors for engines such as the LM2500 or Allison 501 are sufficient for prognostic health monitoring analyses, then transfer of the technology and methods described in this paper may be appropriate. It might even be possible to host this prognostic functionality on existing onboard systems.

Likewise, new platforms provide opportunities for more effective placement/selection of sensors for health monitoring and prognostics. Application of the technology on systems currently in design/acquisition phases would avoid retrofitting costs in the future.

Migration of this technology to new systems or upgrading health monitoring capabilities on legacy systems may be accomplished using the same process that has been used for the Army application discussed in this paper:

- Conduct a thorough engineering and failure modes analysis of the target system to identify maintenance issues and high-frequency/high-cost faults and failures
- Identify sensors that are required to diagnose or predict the faults. Determine if appropriate sensors are already in the field (for legacy systems) or in the system design (for new systems).

- Conduct a cost-benefit analysis to determine the impacts on maintenance costs as well as other less tangible impacts on readiness, logistics footprint, and life-cycle management
- If additional sensors are necessary, design and fabricate a prototype data acquisition system.
- Collect data. Test cell and field data are needed, and also maintenance records. For ships that are already contributing sensor data via the ICAS system, it is possible that a fair amount of data already exists.
- With the collected field data, develop/apply prognostic algorithms in a desktop version of the system. Based on the results, determine if estimated benefits justify possible costs of developing onboard capability, and if funding can be obtained, develop and demonstrate onboard operational prototype.

CONCLUSIONS

This research has achieved the following major accomplishments:

- A high-level architecture or framework for prognostics was developed and described. This helped to communicate logistics requirements and organizational concepts that are fundamental to establishing capabilities for anticipatory logistics.
- A novel, LEAP-Frog regression method was developed to provide more adaptive predictions of future performance from dynamic data. This method was a key element of an invention and patent that was submitted by PNNL on prognostics methods (Patent E-1789). A version of the LEAP-Frog method has been incorporated in the REDI-PRO prototype prognostics system for the U.S. Army.
- The high level LEAP Prognostics architecture and the LEAP-Frog analytic method are applicable to a variety of prognostic problems.

REFERENCES

Greitzer, F.L. 1999. Life Extension Analysis and Prognostics (LEAP) Architectures. In *Laboratory-Directed Research and Development Annual Report, Fiscal Year 1999*. PNNL-13203, Pacific Northwest National Laboratory, pp. 85-88.

Greitzer, F. L., E.J. Stahlman, T.A. Ferryman, W. Wilson, L.J. Kangas, and D.R. Sisk. 1999. "Development of a Framework for Predicting Life of Mechanical Systems: Life Extension Analysis and Prognostics (LEAP)." Presented at SOLE *'99 Symposium*, August 31- September 2, 1999, Las Vegas, Nevada.

ACKNOWLEDGMENT

This work was performed for the U.S. Department of Energy by the Pacific Northwest National Laboratory under contract DE-AC06-76RL01830.

The authors wish to thank Ron Pawlowski and Mario Bagaglio for their thoughtful contributions to this paper.

AUTHOR INFORMATION

FRANK L. GREITZER, PhD
Staff Scientist
Battelle, Pacific Northwest National Laboratory
P.O. Box 999
Richland, WA 99352

Telephone:	(509) 372-4251
Fax:	(509) 372-4913
Email:	frank.greitzer@pnl.gov

THOMAS A. FERRYMAN, PhD
Staff Scientist
Battelle, Pacific Northwest National Laboratory
P.O. Box 999
Richland, WA 99352

Telephone:	(509) 375 3888
Fax:	(509) 375 2604
Email:	tom.ferryman@pnl.gov

BIOGRAPHY

Dr. Frank L. Greitzer

Dr. Greitzer, a Staff Scientist at the Pacific Northwest National Laboratory (PNNL), holds a PhD in mathematical psychology and a BS in mathematics. His professional experience includes over twenty years of advanced technology research and development on performance support systems, human information processing and decision making, and artificial intelligence/expert systems for U.S. Army, Navy, and Air Force applications. Dr. Greitzer manages the REDI-PRO (Real time Engine Diagnostics-Prognostics) project, which is developing a prototype onboard system to diagnose and predict faults in the Army's M1A1 tank gas turbine engine. He also leads the LEAP (Life Extension Analysis and Prognostics) Laboratory-Directed Research and Development project at PNNL, which is developing advanced techniques for predicting remaining useful life of mechanical systems. A member of the Human Factors and Ergonomics Society, Dr. Greitzer is also engaged in the design and development of web-based interactive training systems and user-centered design of human information interfaces.

Dr. Thomas A. Ferryman

Dr. Ferryman, a Staff Scientist at PNNL, holds a PhD in applied statistics, MS in statistics, and a master's degree in business administration. He has over twenty years experience in developing scientific solutions to complex problems, such as: real-time forecasting methods for monitoring the health of complex electro/mechanical systems; aviation safety analysis tools for NASA that may be used by commercial airlines to identify typical flight patterns, atypical flights that warrant safety investigations, operational trends and maintenance trends; and data analysis tools for the U.S. Navy to use in identifying typical and atypical performance of major caliber gun weapon systems on surface ships. Dr. Ferryman also provides systems engineering leadership for over 200 engineers, scientists and technicians in the development, installation and testing of a modification to the AC-130H gunship (a modified transport aircraft) for the U.S. Special Operation Forces. Dr. Ferryman is a member of the American Statistical Association.